

TWO-DIMENSIONAL CORE-COLLAPSE SUPERNOVA MODELS WITH MULTI-DIMENSIONAL TRANSPORT

JOSHUA C. DOLENCE¹, ADAM BURROWS¹, WEIQUN ZHANG²

Draft version March 26, 2014

ABSTRACT

We present new two-dimensional (2D) axisymmetric neutrino radiation/hydrodynamic models of core-collapse supernova (CCSN) cores. We use the CASTRO code, which incorporates truly multi-dimensional, multi-group, flux-limited diffusion (MGFLD) neutrino transport, including all relevant $\mathcal{O}(v/c)$ terms. Our main motivation for carrying out this study is to compare with recent 2D models produced by other groups who have obtained explosions for some progenitor stars and with recent 2D VULCAN results that did not incorporate $\mathcal{O}(v/c)$ terms. We follow the evolution of 12, 15, 20, and 25 solar-mass progenitors to approximately 600 milliseconds after bounce and do not obtain an explosion in any of these models. Though the reason for the qualitative disagreement among the groups engaged in CCSN modeling remains unclear, we speculate that the simplifying “ray-by-ray” approach employed by all other groups may be compromising their results. We show that “ray-by-ray” calculations greatly exaggerate the angular and temporal variations of the neutrino fluxes, which we argue are better captured by our multi-dimensional MGFLD approach. On the other hand, our 2D models also make approximations, making it difficult to draw definitive conclusions concerning the root of the differences between groups. We discuss some of the diagnostics often employed in the analyses of CCSN simulations and highlight the intimate relationship between the various explosion conditions that have been proposed. Finally, we explore the ingredients that may be missing in current calculations that may be important in reproducing the properties of the average CCSNe, should the delayed neutrino-heating mechanism be the correct mechanism of explosion.

Subject headings: (stars:)supernovae: general

1. INTRODUCTION

The mechanism underlying the explosive deaths of massive stars remains poorly understood. Colgate & White (1966) first suggested that neutrino energy deposition plays a central role in powering core-collapse supernovae (CCSNe). Ever since, much of the effort in CCSN theory has focused on building increasingly sophisticated neutrino radiation hydrodynamical models, with the hopes of reproducing the properties of CCSNe, including the kinetic energies, debris morphologies, nucleosynthetic yields, and the remnant mass, spin, and velocity distributions. Despite this effort, the best models still fall short of accounting for any of these properties, much less all of them simultaneously. Perhaps even more alarming, the various groups involved in CCSN modeling often reach qualitatively different conclusions with calculations that are ostensibly quite similar, vis-à-vis whether an explosion even occurs.

In Müller et al. (2012b,a), results are reported from 2D axisymmetric modeling with the VERTEX-CoCoNuT code, which uses a conformally-flat spacetime approximation of general relativity (Müller et al. 2010). They find explosions for 8.1- M_{\odot} , 11.2- M_{\odot} , 15- M_{\odot} , and 27- M_{\odot} progenitors, but, when reported, the explosion energies are ~ 10 times smaller than the canonical 10^{51} erg energy of typical CCSNe. Moreover, they find no explosions for a variety of other progenitors of intermediate masses (Müller et al., YITP Conference, 2013). Recently, Hanke et al. (2013) reported results from a three-dimensional simulation with the PROMETHEUS-VERTEX code of the same 27- M_{\odot} progenitor considered in Müller et al. (2012a) and found no explosion. This negative result has been recapitulated for all other 3D simulations performed recently by this group (Tamborra et al. 2014), despite their having seen explosions in the corresponding 2D simulations.

Suwa et al. (2010) reports an explosion of a 13- M_{\odot} progenitor in a 2D simulation and Takiwaki et al. (2012) finds explosions of an 11.2- M_{\odot} progenitor in both 2D and 3D. These models neglected the heavy lepton neutrinos, which were recently incorporated with an approximate leakage scheme (Takiwaki et al. 2013). In all cases, the ν_e and $\bar{\nu}_e$ transport was computed using the isotropic diffusion source approximation (IDSA) (Liebendörfer et al. 2009), a crude approximation meant to enable multi-D simulations at minimal cost. While interesting, their results are difficult to interpret in the context of the viability of the neutrino mechanism, as the authors acknowledge.

Meanwhile, Bruenn et al. (2013) report results of 2D axisymmetric modeling with their CHIMERA code. They consider 12- M_{\odot} , 15- M_{\odot} , 20- M_{\odot} , and 25- M_{\odot} progenitors from Woosley & Heger (2007) and find explosions in all cases, curiously at almost the same post-bounce time. They also report energies that are somewhat larger than the those reported in Müller et al. (2012b), but that still fall short of the 10^{51} erg mark. Müller et al. (YITP Conference, 2013) recently reported 2D models of the same four progenitors, and found that only the 15- M_{\odot} model exploded.

Importantly, all of the studies discussed above relied on the so-called “ray-by-ray-plus” approximation of neutrino

¹ Department of Astrophysical Sciences, Princeton University, Princeton, NJ 08544; jdolence@astro.princeton.edu, burrows@astro.princeton.edu

² Center for Computational Sciences and Engineering, Lawrence Berkeley National Laboratory, Berkeley, CA 94720; weiqun-zhang@lbl.gov

transport, which replaces the real transport problem with a series of independent spherically-symmetric transport solves. This is a crude approximation that introduces large variations in angle and time in the neutrino fluxes and the associated neutrino energy deposition so crucial for the neutrino-driven mechanism. This simplification has yet to be clearly justified, and may be producing qualitatively incorrect results, particularly in 2D.

The only calculations ever performed which allow for multidimensional transport were the VULCAN/2D results reported in Burrows et al. (2006), Burrows et al. (2007a), Ott et al. (2008), and Brandt et al. (2011), and none of these calculations showed a revival of the stalled shock in 2D by the delayed-neutrino mechanism. The calculations of Ott et al. (2008) and Brandt et al. (2011) were multi-angle as well. However, these calculations were performed without $\mathcal{O}(v/c)$ transport effects (Hubeny & Burrows 2007). We are, therefore, motivated in this paper to perform new 2D multi-group radiation hydrodynamics calculations with a new code with both multi-D transport (avoiding the simplifications of the ray-by-ray approach) and the velocity-dependent terms to determine whether these earlier results were artifacts of the neglect of $\mathcal{O}(v/c)$ terms, and for comparison with the 2D results of other groups. To accomplish this, we have developed the CASTRO radiation hydrodynamics code. CASTRO contains a multi-group flux-limited neutrino transport solver, is time-dependent and multidimensional, and treats three neutrino species (ν_e , $\bar{\nu}_e$, ν_x , where the ν_x species includes the μ and τ neutrinos and their antiparticles), including all relevant $\mathcal{O}(v/c)$ terms. We find that none of our new 2D calculations, employing the same progenitors as Bruenn et al. (2013), explode by the delayed-neutrino mechanism. With this paper, we describe our results and speculate on the reasons for the different outcomes we find. Since all other groups are using the ray-by-ray approach, we suggest that one reason for the different outcomes may be in the handling of multi-D transport. In 2D, the axial sloshing motions, often not seen in 3D (Burrows et al. 2012), may be reinforcing the errors in the ray-by-ray approach and leading to a qualitatively incorrect outcome. In 3D, these axial sloshing effects are often absent, and the ray-by-ray approach may be less anomalous (due to the greater sphericity of the hydrodynamics), so the lack of explosions seen by the Garching group in 3D, when they observe explosions for the same progenitors in 2D, remains puzzling.

2. NUMERICS AND SETUP

We use the CASTRO code to carry out our CCSN simulations (Almgren et al. 2010; Zhang et al. 2011, 2013). CASTRO is a second-order, Eulerian, compressible, Godunov-type, radiation hydrodynamics code that uses block-structured adaptive mesh refinement to simultaneously refine in both space and time. Simulations can be performed in 1D (spherical), 2D (cylindrical), and 3D (Cartesian). The hydrodynamic updates use piecewise-parabolic reconstruction with higher-order limiters to preserve accuracy at smooth extrema, an approximate Riemann solver detailed in Almgren et al. (2010), and incorporates full corner coupling in the directionally unsplit integration. In this work, we make use of the multi-group flux-limited diffusion (MGFLD) neutrino transport solver detailed in Zhang et al. (2013).

Currently, multi-dimensional multi-angle transport is not feasible and will require the exascale. Our comoving frame multi-group flux-limited diffusion (MGFLD) formulation includes $\mathcal{O}(v/c)$ terms that can result (correctly) in significant differences in the dynamic diffusion limit where radiation transport is dominated by motion of the fluid (Castor 2004). The approach used in CASTRO splits the system into three parts: 1) a part that couples the radiation and fluid in a hyperbolic subsystem with a piecewise parabolic method (PPM) with characteristic tracing and full corner coupling (Miller & Colella 2002), 2) another part that is a system of coupled parabolic equations that evolves radiation diffusion over all the groups (along with Y_e in the neutrino case) and source-sink terms, and 3) a final part that performs frequency-space advection. The hyperbolic subsystem is solved explicitly with a high-order Godunov scheme as part of the hydrodynamic component of the algorithm, whereas the parabolic part is solved implicitly with a first-order backward Euler method. Frequency-space advection is performed using a standard approach based on the method of lines. The frequency-space advection has its own CFL condition for stability and, if necessary, subcycling in time is employed in order to satisfy the frequency-space CFL condition. The primary computational expense of the radiation is in the solution of linear systems as part of the iteration over energy groups. We rely on the *hypr* library (Falgout & Yang 2002) for solving these systems on large parallel machines.

CASTRO uses a hybrid parallelization strategy based on MPI + OpenMP using the BoxLib framework (Almgren et al. 2010). The basic strategy is to distribute grids within the AMR hierarchy to computational nodes. This provides a natural coarse-grained approach to distributing the computational work. A dynamic load balancing technique is needed to adjust the load. Although the code supports both a heuristic knapsack algorithm and a space-filling curve algorithm for load balancing, the data-locality properties make the space-filling curve the method of choice for problems with radiation. The main advantages of the CASTRO code are the efficiency due to the use of AMR, the temporal sub-cycling, and the accuracy due to the coupling of radiation force into the Riemann solver.

CASTRO incorporates 1) a complicated set of opacity tables for the various neutrino species that includes weak-magnetism, ion-ion correlation, and ion form-factor corrections, 2) various extant nuclear equations of state (including the Shen [default] and various Lattimer & Swesty equations of state), 3) inelastic neutrino-electron and neutrino-nucleon scattering, and 4) a temporal sub-cycling algorithm that accelerates computation by many factors over traditional codes. CASTRO treats three neutrino species (ν_e , $\bar{\nu}_e$, ν_x , where the ν_x species includes the μ and τ neutrinos and their antiparticles). In this work, for computational expediency, we neglect in the 2D simulations energy-group coupling processes such as inelastic scattering, and justify this choice in Section 4 through the comparison of model shock radius evolutions in 1D, with and without them.

We follow the collapse, bounce, and subsequent evolution of four progenitors in 1D and 2D. The progenitors are nonrotating, solar metallicity models with zero-age main sequence masses of 12- M_\odot , 15- M_\odot , 20- M_\odot , and 25- M_\odot (Woosley & Heger 2007). These are the same progenitors recently considered by Bruenn et al. (2013) and Müller et

al. (YITP conference, 2013). Our numerical grid has 0.5 km resolution in the inner ~ 128 km, 1 km resolution out to ~ 270 km, never worse than 4 km resolution anywhere beneath the shock, and extends out to 5120 km.

3. RESULTS OF TWO-DIMENSIONAL SIMULATIONS

3.1. Global Structure

Figure 1 shows representative snapshots of the radial velocity and entropy from our 12- M_\odot and 25- M_\odot models. At 200 ms after bounce, both models are nearly spherical and show weak convective activity with a characteristic angular scale corresponding to $\ell \sim 10$ (ℓ is spherical harmonic degree) as is clearly seen in the radial velocities. By 400 ms after bounce, both models have become more aspherical, indicating stronger convective activity that is at larger angular scales (smaller ℓ). The snapshots at 600 ms after bounce look qualitatively similar to those at 400 ms after bounce, showing large scale convective motions including relatively low entropy accretion streams and higher entropy buoyant plumes.

Figure 2 shows the evolutions of the average shock radii for our 1D and 2D models. Not surprisingly, our 1D models do not explode, consistent with nearly all spherically symmetric CCSN calculations. In 2D, despite evolving the models to $\gtrsim 600$ ms after bounce, we do not find explosions. Nevertheless, there are interesting differences between 1D and 2D models. Perhaps most importantly, the 2D stalled shock radii are consistently larger by a factor 1.3–2, as found in essentially all 1D-2D comparisons (e.g. Müller et al. 2012b). This is often attributed to aspherical instabilities, namely convection and the standing accretion shock instability (SASI), both increasing the neutrino heating efficiency in the gain region and providing turbulent pressure support (e.g. Herant et al. 1992; Burrows et al. 1995; Janka & Mueller 1996). In our models, we see convective activity develop immediately as the bounce shock moves out. This brief phase of “prompt convection,” seeded by perturbations introduced by our aspherical grid, drives the shock to large radii much faster than in the corresponding 1D models. By ~ 20 ms after bounce, all four 2D models have shock radii of ~ 200 km and these shocks stall ~ 60 ms earlier than the 1D models, which never reach radii much beyond ~ 150 km. Between 300 and 400 ms, the 20- M_\odot and 25- M_\odot models show jumps in the average shock radii, both in 1D and 2D. These jumps correspond to the accretion of the Si/O interface in these models, wherein \dot{M} drops suddenly and the shock responds by moving outward. In 1-D, the jump is only about 10% of the pre-interface shock radius, but the effect is 2–3 times larger in 2D. Apparently and interestingly, the sensitivity to the accretion of shelves increases with shock radius.

The shock radii are only one aspect of the global structure that determines if and when explosions commence. The gain radius, $R_g(\theta)$, is defined as the radius above which the net rate of energy deposition in the gas, integrated over all neutrino energies and summed over species, is positive. The gain region, bounded by the gain and shock radii, contains the mass where neutrino energy deposition, occurring at a rate

$$(\mathcal{H} - \mathcal{C})_{\text{gain}} = \sum_s \int d\Omega \int_{R_g(\theta)}^{R_s(\theta)} r^2 dr \int_{\varepsilon_s^{\min}}^{\varepsilon_s^{\max}} (c\kappa_{s,\varepsilon} E_{s,\varepsilon} - j_{s,\varepsilon}) d\varepsilon, \quad (1)$$

is thought to lead to shock revival. In Eq. 1, $s \in \{\nu_e, \bar{\nu}_e, \nu_x\}$, $R_s(\theta)$ is the angle-dependent shock radius, ε is neutrino energy, κ is the absorption cross section, E is the neutrino energy density, and j is the volume emissivity. As important is the cooling region where accreted material cools and settles onto the proto-neutron star, undermining the pressure support crucial to reviving the stalled shock. We define a net cooling rate analogous to the heating rate above as

$$(\mathcal{H} - \mathcal{C})_{\text{cool}} = \sum_s \int d\Omega \int_{R_{\nu_s}(\theta)}^{R_g(\theta)} r^2 dr \int_{\varepsilon_{\nu_s}^{\min}}^{\varepsilon_{\nu_s}^{\max}} (c\kappa_{s,\varepsilon} E_{s,\varepsilon} - j_{s,\varepsilon}) d\varepsilon, \quad (2)$$

where R_{ν_s} is the angle-, species-, and energy-dependent neutrinosphere radius defined implicitly by

$$\int_{R_{\nu_s}}^{\infty} \kappa dr = \frac{2}{3}. \quad (3)$$

This definition roughly accounts for all the optically-thin cooling in the region, but does not include diffusive flux from beneath the $\tau = 2/3$ surface. We define an effective inner cooling radius as

$$R_{\text{in}} = \frac{\sum_s \sum_{\varepsilon} |w_s| R_{\nu_s}}{\sum_s \sum_{\varepsilon} |w_s|}, \quad (4)$$

where

$$w_s = \int_{R_{\nu_s}}^{\infty} (c\kappa E - j) r^2 dr \quad (5)$$

is an integral of the net neutrino energy deposition along a column (at fixed θ) and should be understood to be angle-, species-, and energy-dependent. This inner cooling radius is related to the neutrinosphere radii of the dominant cooling agents. The region between R_{in} and R_g is cooling rapidly by optically-thin neutrino emission, predominantly through the ν_e and $\bar{\nu}_e$ species, and represents at least one component of the accretion luminosity. Figure 3 shows the evolutions of the shock, gain, and inner radii for our four 2D models. After a short-lived initial transient, the gain and inner radii generally recede in all four models, reflecting the contraction of the inner proto-neutron star.

3.2. Diagnostic Quantities

The evolutions of certain diagnostic quantities have proven useful in distilling insight from the complicated multi-D dynamics of pre-explosion supernova cores. The heating efficiency, ϵ_h , is defined as the ratio of the net heating in the gain region to the sum $L_{\nu_e} + L_{\bar{\nu}_e}$. We also define an analogous cooling efficiency, $\epsilon_c \equiv -(\mathcal{H} - \mathcal{C})_{\text{cool}}/(L_{\nu_e} + L_{\bar{\nu}_e})$. Both efficiencies are shown in Fig. 4. The heating efficiencies initially rise slowly and show broad peaks around 200 ms after bounce. The peak values range from 6% to 9%, monotonically increasing with progenitor mass. At low optical depths, the heating efficiency is equivalent to an effective optical depth, giving a heating rate $(L_{\nu_e} + L_{\bar{\nu}_e})(1 - \exp(-\tau_{\text{eff}})) \approx (L_{\nu_e} + L_{\bar{\nu}_e})\tau_{\text{eff}}$. The basic character of the heating efficiencies can be understood by a simple estimate of τ_{eff} . First, write $\tau_{\text{eff}} \sim \langle \rho \rangle_g \sigma \Delta R$, where $\langle \rho \rangle_g$ is the average density in the gain region, σ is a characteristic neutrino absorption cross section per unit mass, and $\Delta R = \langle R_s \rangle - \langle R_g \rangle$. The average density is easily computed from the gain mass and shock and gain radii. We write the cross section as $\sigma \approx \sigma_0(\varepsilon/\varepsilon_0)^2 \sim \sigma_0(R_0/R_{\text{in}})^2$, where the 0-subscript signifies some characteristic value. This last scaling is justified by the approximately inverse relationship between root-mean-square neutrino energy and inner cooling radius. Finally, we arrive at

$$\tau_{\text{eff}} \sim \frac{3\sigma_0 M_g}{4\pi} \frac{R_s - R_g}{R_s^3 - R_g^3} \left(\frac{R_0}{R_{\text{in}}} \right)^2. \quad (6)$$

This estimate reproduces the heating efficiencies reasonably well, accounting, for example, for the slow decline in the 12- M_\odot model, the nearly constant value for the 15- M_\odot model, and the rise in the latter half of the 20- M_\odot and 25- M_\odot models. The virtue of this estimate is that it shows how the heating efficiencies depend on the global structure of the solutions.

The cooling efficiency ϵ_c represents the fraction of the ν_e and $\bar{\nu}_e$ luminosities arising from optically thin cooling. Since some of the accreted material advects into optically thick regions before cooling appreciably, the net cooling rate $(\mathcal{H} - \mathcal{C})_{\text{cool}}$ entering into the definition of ϵ_c is not equal to the total accretion luminosity, which must be $\approx GM_{\text{pns}}\dot{M}_{\text{pns}}/R_{\text{pns}}$ (pns=proto-neutron star), at least in an integral-averaged sense. Rather, ϵ_c represents a response function to sudden changes in \dot{M} . For example, if the accretion rate drops due to the accretion of a composition interface, the immediate response in the luminosity (delayed somewhat by the advection time between the shock and cooling region) will be to drop fractionally by approximately $\epsilon_c \dot{M}/\dot{M}_0$, where \dot{M}_0 and \dot{M} are the pre- and post-interface accretion rates, respectively. The noticeable shelves accreted in the 12- M_\odot , 20- M_\odot , and 25- M_\odot models confirm this basic picture. Apparently, a low cooling efficiency may make a model prone to explosion when shelves are accreted since, though \dot{M} can drop appreciably, the luminosity responds only weakly, leaving the model closer to (or perhaps beyond) the critical luminosity for explosion (Burrows 2013). Evidently, this effect is not sufficient to drive explosions in our models.

A variety of “explosion conditions” have been proposed in the literature. The most popular considers the ratio of the timescale for advection through the gain region to the heating timescale. Define the advection timescale as $\tau_{\text{adv}} = M_{\text{gain}}/\dot{M}$ and the heating timescale as

$$\tau_{\text{heat}} = \frac{\int_{\text{gain}} (\rho e - \rho e_0) dV}{(\mathcal{H} - \mathcal{C})_{\text{gain}}}, \quad (7)$$

where e_0 is approximately the zero-point of the EOS, i.e. $e_0 = e(\rho, Y_e, T_{\text{min}})$ where T_{min} is the lowest temperature of the tabulated EOS (10^{-2} MeV). The left panel of Fig. 5 shows the evolutions of this ratio for all four models. Surprisingly, the curves all show the same basic structure, peaking around 200 ms after bounce before slowly declining. Superimposed on these slowly varying curves, the 20- M_\odot and 25- M_\odot models show narrow peaks around 300 ms after bounce, corresponding to the accretion of the Si/O interfaces in these models when \dot{M} drops abruptly by a factor of two.

Another interesting dimensionless quantity that arises naturally in discussing critical conditions for explosion (see, e.g., Dolence et al. 2013) is the ratio of heating in the gain region to the accretion power $\mathcal{H}_g R_s/(GM\dot{M})$, shown in the right panel of Fig. 5. Comparing the left and right panels of Fig. 5, it is clear that these dimensionless ratios have very similar behaviors. Writing the heating in the gain region $\mathcal{H}_g = h_g M_g$, where h_g is the specific heating rate and M_g is the mass in the gain region, and noting that the zero-point-subtracted specific internal energy is of order GM/R_s , it is easy to see that $\mathcal{H}_g R_s/(GM\dot{M}) \sim \tau_{\text{adv}}/\tau_{\text{heat}}$, so the similarities between the left and right panels of Fig. 5 are not surprising. Similarly, the “antesonic” condition, proposed by Pejcha & Thompson (2012), is easily related to these other conditions. Apparently, $\tau_{\text{adv}}/\tau_{\text{heat}} \sim \mathcal{H}_g R_s/(GM\dot{M}) \propto c_s^2/v_{\text{esc}}^2$.

While all three ratios are interesting diagnostics and can clearly indicate when explosions are underway, none seem to have a well-determined critical threshold above which explosions are inevitable. In particular, the value of one is not necessarily associated with transitioning to explosion. The critical values likely depend in detail on how these quantities are defined, and may be below or above one, but reasonable definitions seem to at least give critical values of order one. Of course, since we do not obtain explosions, we are unable to determine these critical values for our models and definitions.

4. COMPARISONS WITH ONE-DIMENSIONAL RESULTS

Multidimensional effects seem to be crucial if the neutrino mechanism succeeds in producing explosions. A variety of effects may be important. The post-shock turbulent flow, driven by convective and/or SASI activity, leads to longer dwell times in the gain region on average, exposing the accreted material to more net heating (Murphy & Burrows 2008; Dolence et al. 2013). The turbulent flow itself gives rise to Reynolds stresses that can help support the shock (Murphy et al. 2013). Convection in the proto-neutron star can advect neutrinos closer to their neutrinosphere, shortening their escape time and boosting the diffusive luminosity from the core (Dessart et al. 2006; Mueller & Janka 2014). Even the efficiency of cooling through the accretion luminosity can vary between dimensions. We have already explored the former two effects in parametrized setups (Dolence et al. 2013; Murphy et al. 2013), so here we focus on characterizing how the neutrino luminosities and average energies differ between 1D and 2D models.

Figure 6 shows the evolutions of the ν_e and $\bar{\nu}_e$ luminosities and root-mean-square neutrino energies, as measured in the laboratory frame at 500 km, for all four progenitor models. Some differences between 1D and 2D seem to be model-independent. For example, the rms energies are systematically higher in 1D than in 2D by ~ 5 –10%, consistent with results obtained by other groups (Buras et al. 2006). The luminosities, on the other hand, do not show such a generic difference. Around 100 ms post-bounce, the 20- M_\odot and 25- M_\odot models show deficits of up to $\sim 20\%$ in ν_e and $\bar{\nu}_e$ luminosities in 2D relative to 1D. Beyond ~ 200 ms post bounce, all the 2D models tend to have higher ν_e and $\bar{\nu}_e$ luminosities by $\sim 5\%$, likely an effect of Ledoux convection in the proto-neutron star, driven by the destabilizing lepton gradient (Buras et al. 2006; Dessart et al. 2006).

The idea of a critical luminosity for explosion at a given mass accretion rate was first discussed in Burrows & Goshy (1993). Murphy & Burrows (2008) later showed, in the context of parametrized modeling, that the critical luminosity is lower in 2D than in 1D, a result confirmed by other studies (Hanke et al. 2013; Dolence et al. 2013; Couch 2013). Unfortunately, the critical curves that emerged from these studies are not easily adapted to self-consistent radiation hydrodynamic models. One reason is that the critical luminosity almost certainly depends on quantities other than L_{ν_e} and \dot{M} . For example, it may depend on the proto-neutron star mass and radius and/or the shock radius, as suggested in the discussion of the dimensionless quantities in Section 3.2. It also depends on the structure of the flow beneath the gain layer, which can be quite different between parametrized and self-consistent models. Nevertheless, the lower critical luminosity in 2D very likely remains in self-consistent models.

In Fig. 7, we show the evolutions in the L_{ν_e} – \dot{M} plane for all four progenitors in both 1D and 2D. All of the models show relatively flat curves until late times, indicating that the neutrino luminosities remain roughly constant while the accretion rate drops. It is during this phase that a crossing of critical luminosity curve (and subsequent explosion) seem most likely, though this does not occur in our models. As in Fig. 6, the 2D models tend to show higher ν_e luminosities at late times relative to 1D. Naively, this might suggest that 2D models are easier to explode not only because the critical curve is lower, but also because 2D models tend to have higher luminosities at a given mass accretion rate. Of course, this argument neglects other important differences between 1D and 2D including the higher rms neutrino energies in 1D, for example. Rewriting Eq. 6 for the effective optical depth (or, equivalently, the heating efficiency) with $M_g = \dot{M}\tau_{\text{adv}}$ and $R_0/R_{\text{in}} = \varepsilon/\varepsilon_0$, we arrive at

$$\frac{\tau_{\text{eff}}^{2\text{D}}}{\tau_{\text{eff}}^{1\text{D}}} \sim \left(\frac{\tau_{\text{adv}}^{2\text{D}}}{\tau_{\text{adv}}^{1\text{D}}} \right) \left[\left(\frac{R_s - R_g}{R_s^3 - R_g^3} \right)^{2\text{D}} \left(\frac{R_s^3 - R_g^3}{R_s - R_g} \right)^{1\text{D}} \right] \left(\frac{\varepsilon_{\text{rms}}^{2\text{D}}}{\varepsilon_{\text{rms}}^{1\text{D}}} \right)^2. \quad (8)$$

The first term on the right hand side is the ratio of the advection timescales, which favors 2D. The second term (in square brackets) favors 1D since the gain volume is typically much larger in 2D. As we show in Fig. 6, the last term given by the square of the ratio of rms neutrino energies favors 1D. Evidently, though some effects in 2D tend toward lower heating efficiencies, these are overwhelmed by the different advection timescales, leading to higher heating efficiencies in 2D. Now consider the diagnostic given by the ratio of heating in the gain region to accretion power, which Fig. 5 shows is intimately related to the ratio of advection to heating timescales. It is easy to show that the 2D to 1D ratio of this diagnostic is

$$\left(\frac{HR_s}{GM\dot{M}} \right)^{2\text{D}} \bigg/ \left(\frac{HR_s}{GM\dot{M}} \right)^{1\text{D}} \approx \left(\frac{L_{\nu_e}^{2\text{D}}}{L_{\nu_e}^{1\text{D}}} \right) \left(\frac{\tau_{\text{eff}}^{2\text{D}}}{\tau_{\text{eff}}^{1\text{D}}} \right) \left(\frac{R_s^{2\text{D}}}{R_s^{1\text{D}}} \right). \quad (9)$$

At least after the first 100–200 ms, all of these terms favor 2D, making 2D models much easier to explode than corresponding 1D models.

In the 1D and 2D models discussed thus far, we have ignored inelastic scattering processes. Figure 8 shows results from two 1D simulations of the 15- M_\odot progenitor that differ only in whether inelastic scattering on electrons is included. The most marked difference is in the rms energy of the ν_x species, which is lower with inelastic scattering included by about 5 MeV. In principle, depositing 5 MeV per ν_x neutrino represents a significant source of heating, but this energy deposition occurs mainly at large optical depths for the ν_e and $\bar{\nu}_e$ species, so the net effect is substantially muted. For example, the ν_e and $\bar{\nu}_e$ luminosities are higher by only 3–4% and the rms neutrino energies are higher by only $\sim 0.5\%$ in the model with inelastic scattering. The ν_x luminosity with inelastic scattering adjusts so that the total neutrino luminosity, summed over all species, is nearly identical to the model without inelastic scattering, with fractional differences typically less than 0.1%. Thus, though inelastic scattering on electrons does lead to small changes in the luminosities and average energies of each species, the net effect on the structure as a whole is quite small. For example, the shock radii shown in Fig. 8 are almost identical, differing by at most 2% and typically less than 1%.

5. ANGULAR AND TEMPORAL VARIATIONS IN THE NEUTRINO SECTOR

5.1. A Criticism of “Ray-by-ray”

One of the motivations for carrying out this study is to compare our MGFLD neutrino transport results with those obtained with the *de facto* standard “ray-by-ray” formulation currently used by all other groups carrying out radiation hydrodynamic simulations of core-collapse supernovae. Ott et al. (2008) carried out both MGFLD and full multi-angle transport calculations and found that, though MGFLD is only approximate in the semi-transparent and free streaming regimes, the two methods actually agree quite well in the supernova context. No such comparison with “ray-by-ray” has yet been discussed. Here, we take the first step towards this comparison, which captures one of our main criticisms of the technique.

We consider a snapshot from our 2D simulation of the 12- M_{\odot} model 604 ms after bounce (i.e., at the end of the run). We extract the energy integrated fluxes of the electron neutrinos at 500 km as computed by CASTRO’s MGFLD solver. With this snapshot, we do two more transport calculations, ignoring velocity- and time-dependence and scattering. The first calculation gives the full angle- and energy-dependent specific intensities at discrete latitudes at 500 km. We then integrate over energy and take the first moment to get the fluxes. Finally, we compute the angle- and energy-dependent specific intensities under the “ray-by-ray” assumption, then integrate over energy and take the first moment to recover the “ray-by-ray” fluxes. We describe our techniques for the multi-angle and “ray-by-ray” calculations in Appendix A. The simulation snapshot and the results of these transport calculations are shown in Fig. 9. The MGFLD and full multi-angle transport results agree quite well in the general character of the angular variation, despite the neglect of velocity- and time-dependence and of scattering in the simplified multi-angle calculation. The variation is smooth and has a small amplitude. The “ray-by-ray” results, by contrast, show wild variations that close inspection reveals are highly correlated with structures in the flow. The peak-to-peak variation in the “ray-by-ray” result is nearly an order-of-magnitude larger than for MGFLD and about a factor of four larger than for multi-angle transport. While this may already be a concern, the most serious failing of the “ray-by-ray” approach is that these large variations are tightly correlated with hydrodynamic structures along each ray. The neutrino radiation field is produced by an integral over many sources at depth, an effect which can only be truly captured by doing full multi-angle transport and seems reasonably well approximated by MGFLD, but which a “ray-by-ray” approach can never reliably reproduce. The “ray-by-ray” flux depends only on the profile along a given radial ray, introducing large variations which should be washed away by the integral character of transport and which are highly correlated with the hydrodynamic variations of the flow.

5.2. Angular Variations

As has been discussed previously, the angular variations of the neutrino radiation field are muted relative to variations in the matter sector (Brandt et al. 2011). In Fig. 10, we show the normalized standard deviations of the shock radii and ν_e -fluxes (at 500 km) as a function of time. We define the standard deviation (σ) normalized by the mean (μ) of quantity Q as

$$\frac{\sigma}{\mu} = \frac{1}{\langle Q \rangle} \left(\frac{\int_0^\pi (Q(\theta) - \langle Q(\theta) \rangle)^2 \sin \theta d\theta}{\int_0^\pi \sin \theta d\theta} \right)^{1/2}, \quad (10)$$

where the angle brackets indicate solid-angle averaging. Typically, the fractional angular variations in the shock radii are about an order-of-magnitude larger than the corresponding variations in the fluxes. We also show the cross-correlations of the shock radii and ν_e -fluxes, defined by

$$\text{Corr}(R_s, F_{\nu_e}) = \frac{1}{2} \int_0^\pi \frac{(R_s - \langle R_s \rangle)(F_{\nu_e} - \langle F_{\nu_e} \rangle)}{\sigma_{R_s} \sigma_{F_{\nu_e}}} \sin \theta d\theta, \quad (11)$$

where σ is the standard deviation, computed as above, but without the normalization by the mean. All the models show a positive correlation on average, as indicated by the gray dashed lines in the figure. The 20- M_{\odot} and 25- M_{\odot} models show a period ~ 100 ms long of consistently high correlation. Interestingly, these phases follow immediately after the accretion of the significant Si/O interfaces in these models, where the accretion rates drop by a factor ≈ 2 . Importantly, however, though the shock and ν_e -fluxes are correlated in all models, the amplitude of the ν_e -flux variation is much smaller than one would find with a “ray-by-ray” calculation, so the degree to which these variations couple may be quite different than in “ray-by-ray” calculations (Burrows 2013).

Tamborra et al. (2014) recently reported systematic asymmetries in the net lepton-number emission in several 3D models and dubbed this feature LESA (Lepton-number Emission Self-sustained Asymmetry). To investigate whether we see a similar phenomenon, we decompose the lepton-number flux $F_{\nu_e}^n - F_{\bar{\nu}_e}^n$ of our 12- M_{\odot} model (n indicates number instead of energy flux) into spherical harmonics and focus on the evolution of the normalized dipolar component a_1/a_0 . In the model described in Tamborra et al. (2014), deviations from spherical symmetry in the lepton-number flux are intimately related to corresponding deviations in the shock structure. In the bottom panel of Fig. 11, we show the evolutions of a_1/a_0 for the lepton-number flux and the shock surface. For ease of comparison, we have multiplied a_1/a_0 for the lepton-number flux by five. We see no sign of a strong correlation between the two quantities. Indeed, the normalized cross-correlation as a function of time lag shows a broad feature around zero offset, with a modest peak value of ~ 0.3 (1 indicates perfectly correlated signals) at -21 ms, comparable to the advection timescale. Importantly,

the magnitude of the asymmetry is also about an order-of-magnitude smaller than shown in Tamborra et al. (2014). In short, we find no evidence for LESA in our models, but they are limited to 2D and more work is needed before a final judgment on the existence of LESA can be made. On the other hand, in our 2D simulations we do find that the dipolar component of the $\text{sum } F_{\nu_e} + F_{\bar{\nu}_e}$ (and also the sum of number fluxes) is highly correlated with the dipolar component of the shock, as shown in the top panel of Fig. 11. In this panel, a_1/a_0 of $F_{\nu_e} + F_{\bar{\nu}_e}$ is shown, multiplied by ten for ease of comparison with a_1/a_0 of the shock.

5.3. Temporal Variations

An interesting diagnostic of the temporal variations in the neutrino sector is the power spectrum of the signal expected to be detected from a galactic supernova. Following Lund et al. (2010), we compute the event rate expected for the IceCube detector at 256 observer orientations for our 25- M_\odot model, assumed to be at a distance of 10 kpc. We compute the power spectra of these signals and then solid-angle average these to produce the power spectrum shown in Fig. 12. Comparing with Fig. 5 of Lund et al. (2010), we see remarkable agreement between our power spectrum and the “north hemispheric average” reported in that work. Interestingly, their power spectrum produced by averaging the spectra from each ray, in a manner similar to what we have done, shows significantly more power, particularly at high frequencies. Their hemispheric averaging is meant to mask the inherent error in their “ray-by-ray” transport by mimicking the angular integral character of transport. Evidently, our MGFLD transport naturally produces power spectra that agree quite nicely with their hemisphere-averaged results. By contrast, the “ray-by-ray” technique clearly leads to significant overestimates of the variability in the quantity $L_{\bar{\nu}_e} \epsilon_{\text{rms}}^2$, particularly at high frequencies. This quantity bears directly on the energy deposition in the gain region, and so directly on the viability of the neutrino-driven supernova mechanism.

One interesting aspect of the power spectrum that was not addressed by Lund et al. (2010) (or the follow-on references Lund et al. 2012 and Tamborra et al. 2013) is its exponential shape. The dashed line in Fig. 12 is a fit of the power spectrum above 50 Hz with $P(f) = P_0 \exp(-4\pi f\tau)$, where P_0 and τ are free parameters. The resulting fit seems to describe our results quite well, and appears to be consistent with what others have found. At least one way to produce an exponential power spectrum is with Lorentzian pulses of the form

$$L(t) \propto \frac{\gamma}{(t - t_0)^2 + \gamma^2}, \quad (12)$$

where t_0 is the time of the pulse and γ is the half-width at half-maximum. In Fig. 13, we show a fake signal, produced as a constant plus 500 Lorentzian pulses³ at randomly chosen times and with randomly chosen amplitudes. The power spectrum, shown below the signal, has an exponential dependence on frequency, with an e-folding timescale of $4\pi\gamma$. So, one interpretation of the exponential power spectrum of our modeled IceCube signal is that the signal is comprised of many approximately Lorentzian pulses with a typical full-width at half-maximum timescale, extracted from our fit, of $2\tau = 2.3$ ms. This represents very rapid variability. A simplistic picture of how this timescale may emerge is of blobs cooling as they advect through the cooling region. In this picture, we adopt the simple scaling $L_{\text{blob}} \epsilon_{\text{rms}}^2 \sim T^8$ and estimate the timescale $2\tau \sim (8v_r d \ln T/dr)^{-1}$, with v_r the radial velocity. Using the velocity and temperature gradient in the cooling region shown in Fig. 9, our estimate yields $2\tau \sim 2.5$ ms, in surprisingly good agreement with the measured value of 2.3 ms.

6. SUMMARY AND CONCLUSIONS

Using our new multi-group, multi-dimensional radiation hydrodynamics code CASTRO, which incorporates all terms to $\mathcal{O}(v/c)$ in the transport and does not make the ray-by-ray approximation employed by all other groups now modeling core-collapse supernovae, we have simulated in two spatial dimensions the dynamics of four progenitor massive star models. One goal was to determine, using a different code, whether the outcome of our previous simulations using the VULCAN/2D methodology (Burrows et al. 2006, 2007a; Ott et al. 2008) depended upon the absence of the $\mathcal{O}(v/c)$ terms in VULCAN/2D. We have determined that the results are qualitatively the same and, as when employing VULCAN/2D, we do not see explosions by the neutrino heating mechanism after ~ 600 milliseconds after bounce. Both codes perform two-dimensional transport, though using a multi-group flux-limited (MGFLD) formulation. This conclusion concerning the overall outcome of these models (i.e., explosion in 2D, driven by neutrino heating) is in contrast with the results of Bruenn et al. (2013) and Müller et al. (YITP Conference, 2013), who also do not agree one with the other, but who do obtain neutrino-driven explosions in some or all of their 2D simulations.

One is left to ponder the reasons for these remaining differences in the community of researchers engaged in detailed simulations of the core-collapse supernova phenomenon. We have demonstrated that the ray-by-ray approach does not reproduce the correct angular and temporal neutrino field variations, though no one has yet performed the necessary head-to-head ray-by-ray versus correct transport comparisons (all else kept the same). We speculate, however, that the combination of the ray-by-ray approach with the artificiality of the axial sloshing effects manifest in 2D simulations may be the reason the groups using ray-by-ray obtain explosions in 2D (when they do).

Other reasons could be differences in microphysics, the effects of general relativity (GR, which we do not include, and others approximate), or the different transport algorithms employed (even in 1D). We suggest that differences in microphysics are an unlikely source of problems, since it can be shown that the numerous hydrodynamic, thermal, and radiative feedbacks in the core-collapse problem mute the effects of even large changes in the neutrino-matter cross

³ The number of pulses only effects the normalization, not the e-folding timescale of the exponential.

sections and associated emissivities on the dynamic evolution after collapse. In 1D test calculations we have performed in which the ν_e –neutron absorption cross section was changed by a factor of two (both increased and decreased), the resulting stalled shock radii were the same to within a few percent. The effect of GR is not large, but is touted by some to be important enough to consider (Müller et al. 2012b). We agree, but note that GR seems not to be generally requisite for explosions, as demonstrated by the 2D 27- M_\odot models reported in Müller et al. (2012a) that included GR and in Hanke et al. (2013) that used a monopolar gravity approximation with mock GR corrections, nevertheless transitioning to explosion at nearly the same post-bounce time. Differences in the transport algorithms (apart from the ray-by-ray versus multi-D transport issue) could be to blame, and code-to-code comparisons are called for. This was one early motivation for embarking upon this study with CASTRO—to see whether the outcomes were different from those we obtained using VULCAN/2D. But, more inter-group comparisons, not just intra-group comparisons, are needed.

The fact that the 3D simulations of the Garching group are not exploding when they were in 2D (Hanke et al. 2013) should be a wake-up call to the community to determine the origins of these differences between the various simulation groups and between 2D and 3D. As we have suggested, the use of the ray-by-ray approach is dubious, and since its artificial character is more manifest in 2D we suspect that it is part of the problem. However, this does not explain the current conundrum in 3D—something else may be amiss. It could be that the progenitor models are to blame and a new generation of such models, performed in 3D at the terminal stages of a massive star’s life are needed (Meakin et al. 2011). It could be that rotation, even the modest rotation expected from the pulsar injection constraint (Emmering & Chevalier 1989), could, by the resultant centrifugal support and consequent expected boost in the stalled shock radius, convert duds into explosions. This is the simplest solution, and one is reminded that the exploding model of Marek & Janka (2009) was rotating. Both large-scale and turbulent magnetic fields could play a role, through the associated stress, but also due to enhanced angular momentum transport from the core to the mantle (e.g. Sawai & Yamada 2014). However, without very rapid rotation, which might be associated with the rare hypernovae (Burrows et al. 2007b) that serve as a bridge to the collapsar model of long-soft gamma-ray bursts, there would not seem to be enough extra free energy to power explosion generically. Perturbations of the progenitor cores that collapse have never been properly included into supernova theory, and might be a fruitful line of investigation (Couch & Ott 2013). Such perturbations seed the instabilities long identified with more robust dynamics and the viability of the delayed neutrino mechanism.

Whatever the solution to this recalcitrant problem, advances in the numerical arts seem destined to play a central role. Approximations have been made by all groups to accommodate the limitations of the available computer resources, leaving one to wonder whether such compromises have corrupted the results. One would hope that simple, compelling reasoning, and physical insight could in the end lead to a solution. This has happened before in astrophysics. However, the complexity of the dynamics, the fact that the explosion energy is a small fraction of the available energy, and the circumstance that the central “engine” is shrouded in mystery by the profound opacity of the stellar envelope, and, hence, is itself (almost) inaccessible to direct observation or measurement, may mitigate against a breakthrough unaided by computation.

The authors acknowledge conversations and collaborations with Jeremiah Murphy, Christian Ott, Stan Woosley, Ann Almgren, John Bell, and Louis Howell. The development of the CASTRO code was supported by the Scientific Discovery through Advanced Computing (SciDAC) program of the DOE, under grant number DE-FG02-08ER41544, the NSF under the subaward no. ND201387 to the Joint Institute for Nuclear Astrophysics (JINA, NSF PHY-0822648), and the NSF PetaApps program, under award OCI-0905046 via a subaward no. 44592 from Louisiana State University to Princeton University. The authors employed computational resources provided by the TIGRESS high performance computer center at Princeton University, which is jointly supported by the Princeton Institute for Computational Science and Engineering (PICSciE) and the Princeton University Office of Information Technology; by the National Energy Research Scientific Computing Center (NERSC), which is supported by the Office of Science of the US Department of Energy under contract DE-AC03-76SF00098; and on the Kraken supercomputer, hosted at NICS and provided by the National Science Foundation through the TeraGrid Advanced Support Program under grant number TG-AST100001.

APPENDIX

APPROXIMATE MULTI-ANGLE AND “RAY-BY-RAY” TRANSPORT

For comparisons between the various approaches to neutrino transport, we have computed fluxes based on approximate multi-angle and “ray-by-ray” calculations of the specific intensities. We make three simplifying assumptions: 1) the problem is treated as time-independent, 2) we drop all $\mathcal{O}(v/c)$ terms, and 3) we ignore all scattering processes. Under these assumptions, the transport equation along any ray reduces to

$$\frac{dI_\varepsilon}{ds} = j_\varepsilon - c\kappa_\varepsilon I_\varepsilon, \quad (\text{A1})$$

where I_ε is the specific intensity, s is the path length along the ray, j_ε is the volume emissivity, and κ_ε is the absorption cross section, all defined at neutrino energy ε .

In order to compute the fluxes to facilitate a comparison of multi-angle and “ray-by-ray” with our MGFLD results, we must compute the direction-dependent specific intensities at each point of comparison so that the fluxes can be

computed by taking the first moment. We proceed as follows. First, at each point of comparison, we setup a local spherical coordinate system $(\tilde{r}, \tilde{\theta}, \tilde{\phi})$ with the $\tilde{\theta} = 0$ axis oriented along the local radial direction of the simulation. We then compute the intensities $I(\tilde{\theta}, \tilde{\phi})$ at discrete angles by solving Eq.A1 along the ray passing through our point at angles $\tilde{\theta}$ and $\tilde{\phi}$ to the radial direction.

For the multi-angle calculation, we compute intensities between $\tilde{\theta} = 0$ and $\tilde{\theta} = \tilde{\theta}_{\max}$, where $\tilde{\theta}_{\max}$ is the maximum angular extent of the shock, within which essentially all of the emission/absorption occurs. At each $\tilde{\theta}$, we must also compute intensities as a function of $\tilde{\phi}$, but the axisymmetry of the underlying simulation implies $I(\tilde{\theta}, \tilde{\phi}) = I(\tilde{\theta}, -\tilde{\phi})$, which allows us to consider only $\tilde{\phi}$ in the range $[0, \pi]$.

For the “ray-by-ray” calculation, we must only compute $I(\tilde{\theta})$ since the underlying assumption of a spherical distribution of matter implies that the intensities are symmetric about the local radial direction. This is, of course, the main motivation behind the “ray-by-ray” scheme—there are far fewer intensities to compute and they depend on far less of the simulation data.

REFERENCES

- Almgren, A. S., et al. 2010, *ApJ*, 715, 1221
- Brandt, T. D., Burrows, A., Ott, C. D., & Livne, E. 2011, *ApJ*, 728, 8
- Bruenn, S. W., et al. 2013, *ApJ*, 767, L6
- Buras, R., Janka, H.-T., Rampp, M., & Kifonidis, K. 2006, *A&A*, 457, 281
- Burrows, A. 2013, *Reviews of Modern Physics*, 85, 245
- Burrows, A., Dessart, L., Livne, E., Ott, C. D., & Murphy, J. 2007a, *ApJ*, 664, 416
- . 2007b, *ApJ*, 664, 416
- Burrows, A., Dolence, J. C., & Murphy, J. W. 2012, *ApJ*, 759, 5
- Burrows, A., & Goshy, J. 1993, *ApJ*, 416, L75
- Burrows, A., Hayes, J., & Fryxell, B. A. 1995, *ApJ*, 450, 830
- Burrows, A., Livne, E., Dessart, L., Ott, C. D., & Murphy, J. 2006, *ApJ*, 640, 878
- Castor, J. I. 2004, *Radiation Hydrodynamics*
- Colgate, S. A., & White, R. H. 1966, *ApJ*, 143, 626
- Couch, S. M. 2013, *ApJ*, 775, 35
- Couch, S. M., & Ott, C. D. 2013, *ApJ*, 778, L7
- Dessart, L., Burrows, A., Livne, E., & Ott, C. D. 2006, *ApJ*, 645, 534
- Dolence, J. C., Burrows, A., Murphy, J. W., & Nordhaus, J. 2013, *ApJ*, 765, 110
- Emmering, R. T., & Chevalier, R. A. 1989, *ApJ*, 345, 931
- Falgout, R. D., & Yang, U. M. 2002, in *Computational Science—ICCS 2002* (Springer Berlin Heidelberg), 632–641
- Hanke, F., Müller, B., Wongwathanarat, A., Marek, A., & Janka, H.-T. 2013, *ApJ*, 770, 66
- Herant, M., Benz, W., & Colgate, S. 1992, *ApJ*, 395, 642
- Hubeny, I., & Burrows, A. 2007, *ApJ*, 659, 1458
- Janka, H.-T., & Mueller, E. 1996, *A&A*, 306, 167
- Liebrandt, M., Whitehouse, S. C., & Fischer, T. 2009, *ApJ*, 698, 1174
- Lund, T., Marek, A., Lunardini, C., Janka, H.-T., & Raffelt, G. 2010, *Phys. Rev. D*, 82, 063007
- Lund, T., Wongwathanarat, A., Janka, H.-T., Müller, E., & Raffelt, G. 2012, *Phys. Rev. D*, 86, 105031
- Marek, A., & Janka, H.-T. 2009, *ApJ*, 694, 664
- Meakin, C. A., Sukhbold, T., & Arnett, W. D. 2011, *Ap&SS*, 336, 123
- Miller, G. H., & Colella, P. 2002, *Journal of Computational Physics*, 183, 26
- Mueller, B., & Janka, H.-T. 2014, *ArXiv e-prints*
- Müller, B., Janka, H.-T., & Dimmelmeier, H. 2010, *ApJS*, 189, 104
- Müller, B., Janka, H.-T., & Heger, A. 2012a, *ApJ*, 761, 72
- Müller, B., Janka, H.-T., & Marek, A. 2012b, *ApJ*, 756, 84
- Murphy, J. W., & Burrows, A. 2008, *ApJ*, 688, 1159
- Murphy, J. W., Dolence, J. C., & Burrows, A. 2013, *ApJ*, 771, 52
- Ott, C. D., Burrows, A., Dessart, L., & Livne, E. 2008, *ApJ*, 685, 1069
- Pejcha, O., & Thompson, T. A. 2012, *ApJ*, 746, 106
- Sawai, H., & Yamada, S. 2014, *ArXiv e-prints*
- Suwa, Y., Kotake, K., Takiwaki, T., Whitehouse, S. C., Liebrandt, M., & Sato, K. 2010, *PASJ*, 62, L49
- Takiwaki, T., Kotake, K., & Suwa, Y. 2012, *ApJ*, 749, 98
- . 2013, *ArXiv e-prints*
- Tamborra, I., Hanke, F., Janka, H.-T., Mueller, B., Raffelt, G. G., & Marek, A. 2014, *ArXiv e-prints*
- Tamborra, I., Hanke, F., Müller, B., Janka, H.-T., & Raffelt, G. 2013, *Physical Review Letters*, 111, 121104
- Woosley, S. E., & Heger, A. 2007, *Phys. Rep.*, 442, 269
- Zhang, W., Howell, L., Almgren, A., Burrows, A., & Bell, J. 2011, *ApJS*, 196, 20
- Zhang, W., Howell, L., Almgren, A., Burrows, A., Dolence, J., & Bell, J. 2013, *ApJS*, 204, 7

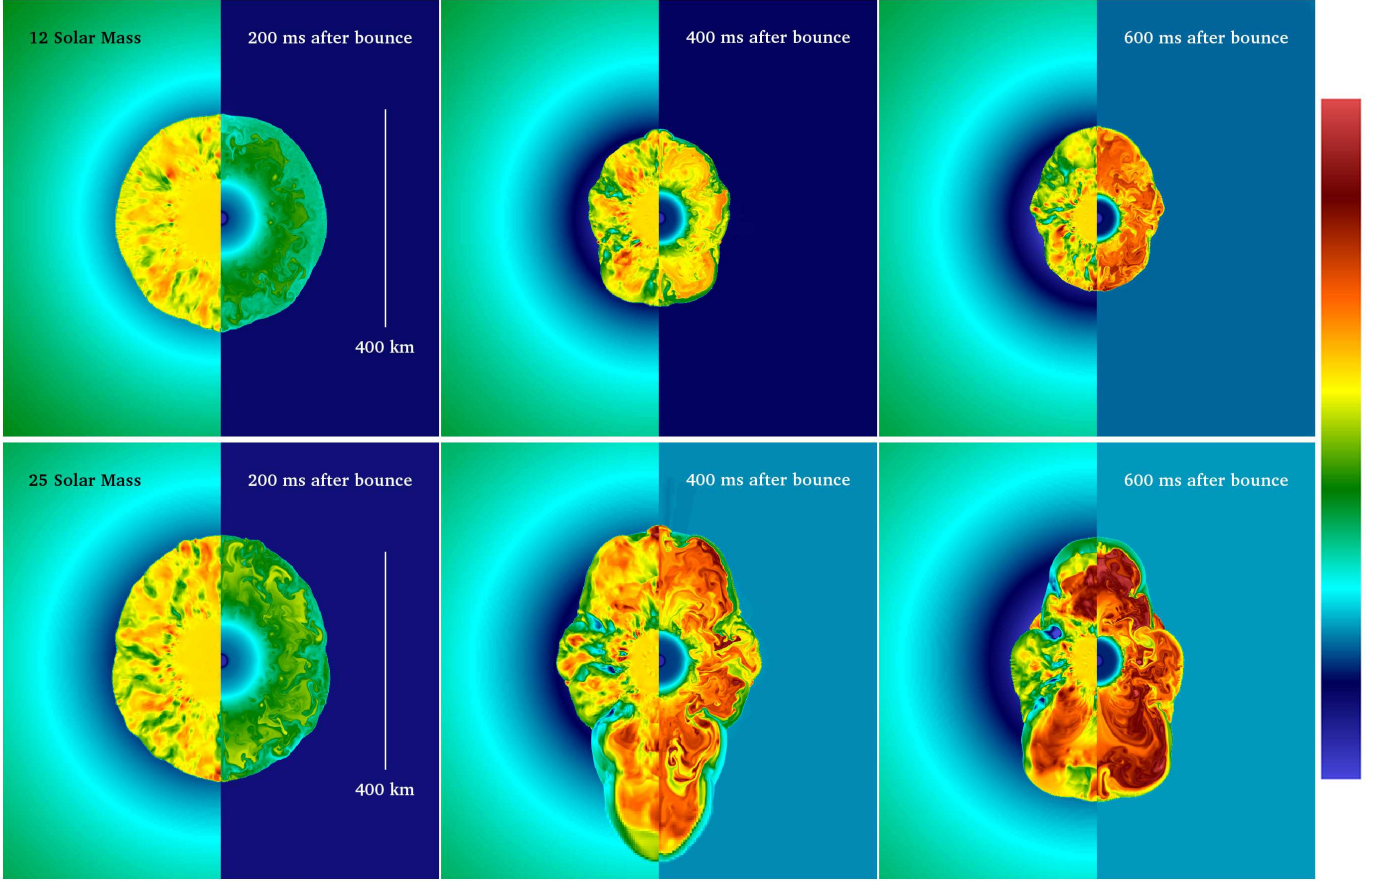


Figure 1. Snapshots of the radial velocity (left half of each image) and entropy (right half of each image) for the 12- M_{\odot} (top three images) and 25- M_{\odot} (bottom three images) models at 200, 400, and 600 ms after bounce. The colorbar, shown at right, ranges from $[-6 \times 10^4, 4 \times 10^4] \text{ km s}^{-1}$ for the radial velocity and from $[0, 25] k_B \text{ baryon}^{-1}$ for the entropy. Convective activity can be clearly seen as relatively low entropy accretion streams and upwelling higher entropy plumes.




Figure 2. Evolutions of the average shock radii for our 1D and 2D simulations. The 2D models diverge from their 1D counterparts almost immediately, likely due to a brief but vigorous phase of prompt convection seeded by our aspherical grid. The 2D average shock radii then remain larger throughout the evolution. Around 300 ms after bounce, the 20- M_{\odot} and 25- M_{\odot} models show jumps due to the accretion of the Si/O interface and the jump is fractionally larger in 2D than in 1D.




Figure 3. Evolutions of the average shock, gain, and inner radii, as defined in Sec. 3.1. After the initial expansions of the bounce shocks, the gain regions form and the gain radii move inward. This is accompanied by the inward migration of the inner cooling radius. Though the shock radii respond to the accretion of the Si/O interface in the 20- M_{\odot} and 25- M_{\odot} models, the gain and inner cooling radius do not.




Figure 4. Evolutions of the heating and cooling efficiencies as defined in Sec. 3.2. The heating efficiencies rise slowly into a broad peak around 200 ms in all models, whereas the cooling efficiencies show minima around 100 ms before growing for the remainder of the simulations.




Figure 5. The left panel shows the evolutions of the ratio of advection to heating time scales. All four models show a broad peak around 200 ms after bounce, around the same time the heating efficiencies are largest. The narrow peaks around 300 ms for the 20- M_{\odot} and 25- M_{\odot} models correspond to the accretion of the Si/O interface when \dot{M} drops by a factor of two. The right panel shows the evolutions of the dimensionless ratios of heating to accretion power. Aside from the overall normalization, these curves have a striking resemblance to those shown in the left panel, demonstrating that the ratio of advection to heating timescales is intimately related to the ratio of heating to accretion power.




Figure 6. Evolution of the neutrino luminosities (solid lines, left ordinate) and root-mean-square neutrino energies (dashed lines, right ordinate) for our 1D (thin lines) and 2D (thick lines) simulations of all four progenitors. The 1D models tend to have slightly higher rms neutrino energies, but, at least at later times, slightly lower luminosities.




Figure 7. Trajectories in the L_{ν_e} - \dot{M} plane for all four progenitors in 1D and 2D. Interestingly, all four models show a roughly horizontal track, indicating that the luminosities are roughly constant while the accretion rates drop. The late time luminosity drops of the 20- M_{\odot} and 25- M_{\odot} models are somewhat delayed relative to the drops in the accretion rates associated with the Si/O interfaces. Evidently, since none of our models explode, the critical curve must be to the left of the data in each plot.




Figure 8. Comparison of 1D simulations of the 15- M_\odot progenitor, with and without inelastic scattering on electrons. The largest change is in the rms energy of the ν_x species, which is lowered by ~ 5 MeV with inelastic scattering included. The shock radii, shown in the bottom panel, are almost identical, demonstrating that the effects of inelastic scattering on electrons is quite small in terms of changes in the global structure of the flow.




Figure 9. Comparison of the transport schemes as discussed in Sec. 5. The pseudocolor map shows $\rho T^6 \exp(-\rho/10^{11})$, roughly representative of the local cooling rate. The lines show the normalized ν_e -fluxes measured in the laboratory frame at 500 km. The normalization scales the fluxes so that their solid-angle average is 100 on the plot. Angular variations in the hydrodynamic sector are clearly, and unphysically, represented in the ν_e -fluxes for the “ray-by-ray” scheme, while the MGFLD and multi-angle fluxes show much smoother variations characteristic of the integral nature of transport.

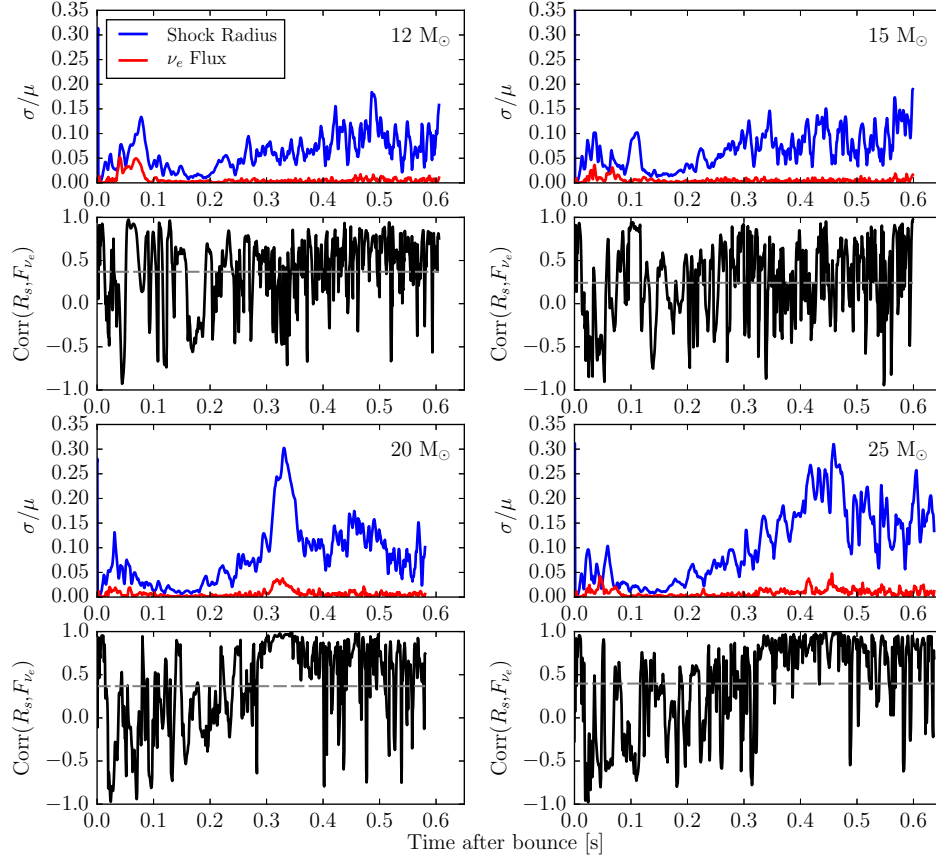


Figure 10. For all four progenitor models, the top panel shows the standard deviations of the shock radius and energy integrated ν_e flux at 500 km, normalized by their respective means. The shock radius shows a fractional variation about an order-of-magnitude larger than the ν_e -fluxes. The bottom panels show the cross-correlations of these quantities, as defined in Sec. 5.2. The dashed gray lines show the value of the time-averaged cross-correlations. See text in Section 5.2 for a discussion.




Figure 11. The bottom panel shows the evolutions of the normalized dipolar component a_1/a_0 of the lepton-number flux $F_{\nu_e}^n - F_{\bar{\nu}_e}^n$ and shock radius, the former multiplied by five for ease of comparison. The signals appear only weakly correlated, which is confirmed quantitatively by computing their cross-correlation. The amplitude of the lepton-number flux asymmetry is also quite small, about an order of magnitude smaller than reported in Tamborra et al. (2014). Taken together, these results seem at odds with the expectations of the LESA model described in Tamborra et al. (2014). The top panel shows the evolutions of the normalized dipolar component of $F_{\nu_e} + F_{\bar{\nu}_e}$ along with the shock, the former multiplied by ten. These quantities are clearly much more correlated than the quantities in the bottom panel.




Figure 12. Power spectrum of the simulated IceCube detection signal, as discussed in Sec. 5.3. The gray dashed line shows a best-fit exponential. The normalization and slope of our power spectrum seems to agree quite well with the results of Lund et al. (2010), who had to compute hemispheric-averages of their “ray-by-ray” data to approximate the effects of multi-D transport.




Figure 13. The top panel shows a fake IceCube signal, produced as a sum of a constant and 500 randomly placed Lorentzian pulses. The bottom panel shows the power spectrum of this fake signal, with the dashed gray line showing the expected exponential fall off for the pulse width adopted in the fake signal, $\exp(-4\pi f\gamma)$, where γ is the half-width at half-maximum, as discussed in Sec. 5.3.